

WHAT IS CLAIMED IS:

1 1. A channel adapter comprising:
2 a host interface, the host interface operatively connected to a memory by a local
3 bus, the memory containing at least one completion queue and at least one event
4 queue;
5 a link interface, the link interface operatively connected to a network;
6 a packet processing engine, the packet processing engine moving data between
7 the host interface and the link interface;
8 an address translation engine, the address translation engine translating a virtual
9 address into a physical address of a translation protection table in the memory; and
10 a completion queue engine, the completion queue engine processing completion
11 requests from the packet processing engine by writing the appropriate at least one of
12 the at least one completion queue and at least one event queue,
13 wherein the packet processing engine is not impacted by any address translation
14 functionality, completion queue accesses, or event queue accesses thereby significantly
15 enhancing the performance of a channel adapter.

1 2. The adapter according to claim 1, wherein the network comprises one of an
2 Infiniband network and a Next Generation Input Output (NGIO) network.

3 3. The adapter according to claim 1, wherein the host interface comprises at
4 least one of a Peripheral Component Interconnect (PCI) interface, PCI-X interface, and
5 a H-16 interface.

4. The adapter according to claim 1, wherein the packet processing engine converts the data from the host interface into packets to be sent to the link interface, and converts packets from the link interface into data to be sent to the host interface.

5. The adapter according to claim 1, wherein the completion queue engine comprises:

at least one local bus register, each at least one local bus register being programmable by at least one of an operating system and an application over the local bus;

at least one context memory, the at least context memory storing context values for each at least one completion queue; and

a finite state machine, the finite state machine processing the completion requests from the packet processing engine and requests generated from the at least one local bus register.

6. The adapter according to claim 5, wherein the at least one local bus register comprises at least one of a completion queue base register, a completion queue entries register, a completion queue protection domain register, a completion queue interrupt enable register, a completion queue number register, and a completion queue control register, the contents of the completion queue base register, completion queue entries register, completion queue protection domain register, and completion queue interrupt enable register being stored in the context memory.

1 7. The adapter according to claim 6, wherein the completion queue engine
2 further comprises at least one first working register, the at least one first working
3 register being loaded with the contents of one of the completion queue base register,
4 completion queue entries register, completion queue protection domain register, and
5 completion queue interrupt enable register stored in the context memory while the finite
6 state machine performs operations on the at least one context queue.

1 8. The adapter according to claim 7, wherein the contents loaded in the at least
2 one first working register is stored back into the context memory after completing the
3 operations.

1 9. The adapter according to claim 7, wherein the contents loaded in the at least
2 one first working register is updated during the performing of operations, the updated
3 contents being stored back into the context memory after completing the operations.

1 10. The adapter according to claim 6, wherein the completion queue control
2 register contains an opcode, the opcode being used by at least one of the operating
3 system and the application to one of enable and disable the at least one completion
4 queue.

1 11. The adapter according to claim 10, wherein the at least one local bus
2 register comprises at least one completion queue host entry register, the values in the
3 at least one completion queue host entry register being added to the at least one

completion queue indicated by the completion queue number register based on the opcode.

12. The adapter according to claim 5, wherein the at least one local bus register comprises at least one of an event queue base register, an event queue entries register, an event queue protection domain register, an event queue interrupt enable register, and an event queue control register.

13. The adapter according to claim 12, wherein the completion queue engine further comprises at least one second working register, the at least one second working register being loaded with the contents of one of the event queue base register, event queue entries register, event queue protection domain register, and event queue interrupt enable register while the finite state machine performs operations on the at least one event queue.

14. The adapter according to claim 13, wherein the event queue control register generates a request to update the values of the event queue base register, event queue entries register, and event queue protection domain register when the event queue control register is written to by at least one of the operating system and the application.

15. The adapter according to claim 5, wherein the at least one local bus register comprises at least one completion queue doorbell register, the at least one completion

queue doorbell register allowing at least one of the operating system and the application to enable automatic event generation for at least one completion queue.

16. The adapter according to claim 1, wherein the address translation engine comprises an inbound request processor and a request completion processor, the inbound request processor receiving a request for address translation of the virtual address, the request completion processor sending a physical address of the memory associated with the virtual address and retrieved from the translation protection table to at least one of the packet processing engine and the completion queue engine in response to the request.

17. The adapter according to claim 16, wherein the inbound request processor comprises:

at least one switching device, the at least one switching device receiving the request for address translation from at least one of the packet processing engine, the completion queue engine, and the request completion processor;

at least one request register bank, each at least one request register bank comprising a request register, and a data register, the request register storing the request, the data register storing data received from at least one of the packet processing engine, the completion queue engine, and the request completion processor related to the request;

arbitration logic, the arbitration logic selecting between all outstanding requests and outputting a protection index of one of the requests, the arbitration logic associating a tag value with each request;

a local bus interface, the local bus interface receiving the protection index, checking whether the protection index is out of bounds, and adding the protection index with a base address of the translation protection table generating the physical address to the translation protection table; and

a request processor, the request processor sending errors related to the request to the request completion processor, the request processor sending a request and the physical address to the host interface to read the translation protection table.

18. The adapter according to claim 17, wherein the local bus interface comprises a size register and at least one base address register, the size register and at least one base address register being programmable over the local bus by at least one of the operating system and the application, the contents of the size register being compared to the protection index to determine whether the protection index is out of bounds, the base address register containing the base address of the translation protection table.

19. The adapter according to claim 17, wherein the data stored by the data register comprises one of a virtual address, a key, the protection index, and a protection domain.

20. The adapter according to claim 16, wherein the request completion processor comprises a decoder, at least one set of receive data buffers, arbitration logic, at least one staging register, permission checking logic, protection index calculation logic, and bounds checking logic.

1 21. The adapter according to claim 20, wherein each at least one set of receive data
2 buffers comprises a read complete register and at least one data buffer, the read complete
3 register signaling a request to the arbitration logic once all at least one data buffer for the
4 set have been filled.

1 22. The adapter according to claim 20, wherein the decoder, at least one set of
2 receive data buffers, and arbitration logic operate at a different clock speed than the at least
3 one staging register, permission checking logic, protection index calculation logic, and
4 bounds checking logic.

1 23. The adapter according to claim 19, further comprising a first valid register
2 associated with the permission checking logic, a second valid register associated with the
3 protection index calculation logic, and a third valid register associated with the bounds
4 checking logic, each valid register containing an indication that the processing for the
5 associated logic may begin.

1 24. A method for enhanced channel adapter performance comprising:
2 receiving a virtual address from a requestor, the virtual address requiring translation
3 to a physical address to a memory;
4 accessing a translation protection table using the virtual address to retrieve at least
5 one data, one at least one data containing a first address;
6 checking the validity of the at least one data;

determining if a second access to the translation protection table is required based on the at least one data;

accessing the translation protection table using the first address to retrieve a second address if required; and

using one of the first address and a physical address to the memory contained at the second address to access the memory,

wherein the packet processing engine is not impacted by any address translation functionality, completion queue accesses, or event queue accesses thereby significantly enhancing the performance of a channel adapter.

25. The method according to claim 24, further comprising adding the virtual address to a base address of the translation protection table to get a physical address to the translation protection table.

26. The method according to claim 24, wherein the at least one data comprises at least one of window entries, region entries, and translation entries, the virtual address accessing at least one of window entries and region entries when part of a key request, the virtual address accessing translation entries when part of a protection index request, each translation entry containing a physical address to the memory.

27. The method according to claim 26, wherein each window entry comprises a first entry containing a key, access rights, and protection information, a second entry containing the first address, and a third entry containing a third address.

1 28. The method according to claim 27, wherein the checking the validity comprises
2 at least one of checking the access rights, comparing a protection key in the virtual address
3 with the key in the first entry, and checking whether the virtual address is within the bounds
4 of the translation protection table.

1 29. The method according to claim 27, further comprising adding the virtual address
2 to the first address to generate the second address.

1 30. The method according to claim 26, wherein region entries comprise a first entry
2 containing a key, access rights, a length of the number of associated translation entries
3 following the first entry, and protection information, a second entry containing the first
4 address, and a third entry containing a translation entry.

1 31. The method according to claim 30, wherein the checking the validity comprises
2 at least one of checking the access rights, comparing a protection key in the virtual address
3 with the key in the first entry, and checking whether the virtual address is within the bounds
4 of the translation protection table.

1 32. The method according to claim 30, further comprising adding the virtual address
2 to the first address to generate the second address.

1 33. A method for enhanced channel adapter performance comprising:

2 completing at least one work task by the channel adapter;
3 storing a completion status for each completed at least one work task in an entry in
4 one at least one completion queue by a completion queue engine, each entry in each at
5 least one completion queue being capable of storing multiple completion statuses;
6 receiving, at the completion queue engine, a completion request associated with at
7 least one work task from a packet processing engine;
8 retrieving the completion status associated with the completion request from the one
9 at least one completion queue by the completion queue engine; and
10 sending the completion status to the packet processing engine.

1 34. The method according to claim 33, further comprising storing notification of the
2 storing of any completion status in one at least one completion queue in an entry in an event
3 queue, the event queue containing an entry for each at least one completion queue.

1 35. The method according to claim 33, further comprising programming at least one
2 register with at least one of a base address, a number of entries, a protection domain, an
3 interrupt enable, and a queue number for one at least one completion queue.

1 36. The method according to claim 35, wherein the programming is performed by
2 at least one of an operating system and an application.

1 37. The method according to claim 34, further comprising programming at least one
2 register with at least one of a base address, a number of entries, a protection domain, and
3 an interrupt enable for the event queue.

1 38. The method according to claim 37, wherein the programming is performed by
2 at least one of an operating system and an application.

1 39. The method according to claim 33, further comprising programming at least one
2 host register with data, the data being written into one at least one completion queue, the
3 programming being performed by at least one of an operating system and an application.

1 40. The method according to claim 34, further comprising programming at least one
2 doorbell register, the at least one doorbell register programmable with at least one of a
3 memory mapped doorbell address and doorbell data, the doorbell data in each at least one
4 doorbell register causing one associated at least one completion queue to generate an entry
5 in the event queue when at least one completion status is stored in the one associated at
6 least one completion queue.

1 41. The method according to claim 40, wherein the programming is performed by
2 at least one of an operating system and an application by directly writing to the at least one
3 doorbell register.

42. The method according to claim 40, wherein each doorbell address is a doorbell stride from each other, a doorbell stride being a relative offset of each doorbell address.

43. The method according to claim 40, wherein the doorbell stride is 4K address locations.

44. The method according to claim 40, wherein the doorbell stride is 16K address locations.

45. A computing device with enhanced channel adapter performance comprising:
a memory, the memory containing at least one translation protection table, at least one completion queue, at least one event queue, and at least one of a data buffer and a one work queue; and

at channel adapter, the channel adapter comprising:

a packet processing engine, the packet processing engine moving data between a first interface and a second interface;

an address translation engine, the address translation engine translating a virtual address into a physical address of the at least one translation protection table in the memory; and

a completion queue engine, the completion queue engine processing completion requests from the packet processing engine by writing the appropriate at least one of the at least one completion queue and the at least one event queue, and

wherein the packet processing engine is not impacted by any address translation functionality, completion queue accesses, or event queue accesses thereby significantly enhancing the performance of the channel adapter.

46. The device according to claim 45, wherein the first interface comprises a host interface, the host interface operatively connected between the channel adapter and the memory by a local bus.

47. The device according to claim 46, wherein the completion queue engine comprises at least one local bus register, each at least one local bus register being programmable by at least one of an operating system and an application over the local bus.

48. The device according to claim 45, wherein the second interface comprises a link interface, the link interface operatively connected between the channel adapter and a network.

49. The device according to claim 48, wherein the network comprises one of an Infiniband network and a Next Generation Input Output (NGIO) network.